


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

 +"generation number" +update "version code" "version number"

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Published before September 2002

Terms used [generation number](#) [update](#) [version code](#) [version number](#)

Found 90 of 132,417

Sort results by

 Save results to a Binder

 Try an Advanced Search

Display results

 Search Tips

 Try this search in [The ACM Guide](#)
 Open results in a new window

Results 1 - 20 of 90

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [next](#)Relevance scale **1 Quorum consensus in nested-transaction systems**

Kenneth J. Goldman, Nancy Lynch

December 1994 **ACM Transactions on Database Systems (TODS)**, Volume 19 Issue 4**Publisher:** ACM PressFull text available: pdf(3.45 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Gifford's Quorum Consensus algorithm for data replication is studied in the context of nested transactions and transaction failures (aborts), and a fully developed reconfiguration strategy is presented. A formal description of the algorithm is presented using the Input/Output automaton model for nested-transaction systems due to Lynch and Merritt. In this description, the algorithm itself is described in terms of nested transactions. The formal description is used to construct a complete pr ...

Keywords: I/O automata, concurrency control, data replication, hierarchical proofs, nested transactions, quorum consensus

**2 A compiler-directed distributed shared memory system**

Tzi-cker Chiueh, Manish Verma

July 1995 **Proceedings of the 9th international conference on Supercomputing****Publisher:** ACM PressFull text available: pdf(1.22 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**3 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data****consistency between database and file systems**

Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02****Publisher:** ACM PressFull text available: pdf(1.44 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management

systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

4 A real-time garbage collector based on the lifetimes of objects 

 Henry Lieberman, Carl Hewitt

June 1983 **Communications of the ACM**, Volume 26 Issue 6

Publisher: ACM Press

Full text available:  pdf(1.37 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In previous heap storage systems, the cost of creating objects and garbage collection is independent of the lifetime of the object. Since objects with short lifetimes account for a large portion of storage use, it is worth optimizing a garbage collector to reclaim storage for these objects more quickly. The garbage collector should spend proportionately less effort reclaiming objects with longer lifetimes. We present a garbage collection algorithm that (1) makes storage for short-li ...

Keywords: LISP, algorithms, languages, lisp, object-oriented programming, parallel processing, performance, real-time garbage collection, reference counting, virtual memory

5 A simulation study of replication control protocols using volatile witnesses 

Perry K. Sloope, Jehan-François Pâris, Darrell D. E. Long

April 1992 **Proceedings of the 25th annual symposium on Simulation ANSS '92**

Publisher: IEEE Computer Society Press

Full text available:  pdf(1.05 MB)

Additional Information: [full citation](#), [references](#), [index terms](#)

Keywords: discrete event simulation, replicated data management, replication control protocols, voting, witnesses

6 BASE: using abstraction to improve fault tolerance 

 Rodrigo Rodrigues, Miguel Castro, Barbara Liskov

October 2001 **ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles SOSP '01**, Volume 35 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.47 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Software errors are a major cause of outages and they are increasingly exploited in malicious attacks. Byzantine fault tolerance allows replicated systems to mask some software errors but it is expensive to deploy. This paper describes a replication technique, BASE, which uses abstraction to reduce the cost of Byzantine fault tolerance and to improve its ability to mask software errors. BASE reduces cost because it enables reuse of off-the-shelf service implementations. It improves availability ...

7 An updated survey of GA-based multiobjective optimization techniques 

Carlos A. Coello

6 June 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 2

Publisher: ACM Press

Full text available: [pdf\(250.77 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

After using evolutionary techniques for single-objective optimization during more than two decades, the incorporation of more than one objective in the fitness function has finally become a popular area of research. As a consequence, many new evolutionary-based approaches and variations of existing techniques have recently been published in the technical literature. The purpose of this paper is to summarize and organize the information on these current approaches, emphasizing the importanc ...

Keywords: artificial intelligence, genetic algorithms, multicriteria optimization, multiobjective optimization, vector optimization

8 Multiview access protocols for large-scale replication

6 Xiangning Liu, Abdelsalam Helal, Weimin Du

6 June 1998 **ACM Transactions on Database Systems (TODS)**, Volume 23 Issue 2

Publisher: ACM Press

Full text available: [pdf\(365.98 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The article proposes a scalable protocol for replication management in large-scale replicated systems. The protocol organizes sites and data replicas into a tree-structured, hierarchical cluster architecture. The basic idea of the protocol is to accomplish the complex task of updating replicated data with a very large number of replicas by a set of related but independently committed transactions. Each transaction is responsible for updating replicas in exactly one cluster and invoking add ...

Keywords: data replication, large-scale systems, multiview access

9 Extended ephemeral logging: log storage management for applications with long lived transactions

6 John S. Keen, William J. Dally

6 March 1997 **ACM Transactions on Database Systems (TODS)**, Volume 22 Issue 1

Publisher: ACM Press

Full text available: [pdf\(566.34 KB\)](#) Additional Information: [full citation](#), [references](#), [index terms](#), [review](#)

Keywords: OLTP, disk management, logging, long transactions

10 XEL: extended ephemeral logging for log storage management

6 John S. Keen, William J. Dally

6 November 1994 **Proceedings of the third international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: [pdf\(987.81 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Extended ephemeral logging (XEL) is a more general variation of the ephemeral logging (EL) technique for managing a log of database activity on disk; it does not require a timestamp to be maintained with each object in the database. XEL does not require periodic checkpoints and does not abort lengthy transactions as frequently as traditional firewall logging for the same amount of disk space. Therefore, it is well suited for

concurrent databases and applications which have a wide distributi ...

11 Distributed object implementations for interactive applications

Vijaykumar Krishnaswamy, Ivan B. Ganev, Jaideep M. Dharap, Mustaque Ahamed
April 2000 **IFIP/ACM International Conference on Distributed systems platforms**

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(175.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

As computers become pervasive in the home and community and homes become better connected, new applications will be deployed over the Internet. Interactive Distributed Applications involve users in multiple locations, across a wide area network, who interact and cooperate by manipulating shared objects. A timely response to user actions, which can potentially update the state of the objects, is an important requirement of interactive applications. Because of the inherent heterogeneity of the ...

12 The iMAX-432 object filing system

 Fred J. Pollack, Kevin C. Kahn, Roy M. Wilkinson
December 1981 **Proceedings of the eighth ACM symposium on Operating systems principles**

Publisher: ACM Press

Full text available:  pdf(1.08 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

iMAX is the operating system for Intel's iAPX-432 computer system. The iAPX-4321 is an object-oriented multiprocessor architecture that supports capability-based addressing.

The object filing system is that part of iMAX that implements a permanent reliable object store. In this paper we describe the key elements of the iMAX object filing system design. We first contrast the concept of an object filing system with that of a conventional file system. We then de ...

13 Towards an efficient management of objects in a distributed environment

 A. El Habbash, J. Grimson, C. Horn
July 1990 **Proceedings of the second international symposium on Databases in parallel and distributed systems**

Publisher: ACM Press

Full text available:  pdf(1.01 MB) Additional Information: [full citation](#), [references](#), [index terms](#)

14 Computer security: Grammar based off line generation of disposable credit card numbers

 Abhishek Singh, Andre L. M. dos Santos
March 2002 **Proceedings of the 2002 ACM symposium on Applied computing**

Publisher: ACM Press

Full text available:  pdf(723.46 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Context free grammars present the desirable cryptographic property that it is easy to generate and validate strings from a given grammar, however it is hard to identify a grammar given only the strings generated by it. The algorithm used in the authentication protocol proposed in this paper makes use of context free grammars. This authentication protocol is a perfect candidate for the offline generation and validation of a disposable credit card number. The proposed protocol can be used alone an ...

15 A methodology for implementing highly concurrent data objects

 Maurice Herlihy
November 1993 **ACM Transactions on Programming Languages and Systems**

(TOPLAS), Volume 15 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.60 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A concurrent object is a data structure shared by concurrent processes. Conventional techniques for implementing concurrent objects typically rely on critical sections; ensuring that only one process at a time can operate on the object. Nevertheless, critical sections are poorly suited for asynchronous systems: if one process is halted or delayed in a critical section, other, nonfaulty processes will be unable to progress. By contrast, a concurrent object i ...

16 A graph-theoretic model for optimizing queries involving methods 

Chiang Lee, Chi-Sheng Shih, Yaw-Huei Chen

April 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**,

Volume 9 Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(266.86 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Traditional algorithms for optimizing the execution order of joins are no more valid when selections and projections involve methods and become very expensive operations. Selections and projections could be even more costly than joins such that they are pulled above joins, rather than pushed down in a query tree. In this paper, we take a fundamental look at how to approach query optimization from a top-down design perspective, rather than trying to force one model to fit into another. We present ...

Keywords: Graph model, Method query, Object-oriented databases, Query optimization, Spanning tree

17 Hierarchical data management 

 Jan M. Engel

September 1976 **Proceedings of the eighth international conference on APL**

Publisher: ACM Press

Full text available:  pdf(1.13 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

An APL program has been developed for storing and maintaining relevant information about a group of persons linked together by a structured hierarchy. Functions that create, modify, and develop useful output from the data set are described with reference to a working example which further illustrates the approach used.

18 Locating and accessing data repositories with WebSemantics 

George A. Mihaila, Louisa Raschid, Anthony Tomasic

August 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 1

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(130.61 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Many collections of scientific data in particular disciplines are available today on the World Wide Web. Most of these data sources are compliant with some standard for interoperable access. In addition, sources may support a common semantics, i.e., a shared meaning for the data types and their domains. However, sharing data among a global community of users is still difficult because of the following reasons: (i) data providers need a mechanism for describing and publishing available sources of ...

Keywords: Data discovery, Data integration, Mediators, Query languages, World Wide Web, XML

19 How to reuse a "write - once" memory (Preliminary Version)

Ronald L. Rivest, Adi Shamir

May 1982 **Proceedings of the fourteenth annual ACM symposium on Theory of computing**

Publisher: ACM Press

Full text available: [pdf\(679.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Storage media such as digital optical disks, PROMS, or paper tape consist of a number of -“write-once-&-rdquo; bit positions (wits); each wit initially contains a -“0-&-rdquo; that may later be irreversibly overwritten with a -“1-&-rdquo;. We demonstrate that such -“write-once memories-&-rdquo; (woms) can be -“rewritten-&-rdquo; to a surprising degree. For example, only 3 wits suffice to represent any 2-bit ...

20 Low-loss TCP/IP header compression for wireless networks

Mikael Degermark, Mathias Engan, Björn Nordgren, Stephen Pink
October 1997 **Wireless Networks**, Volume 3 Issue 5

Publisher: Kluwer Academic Publishers

Full text available: [pdf\(534.08 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Wireless is becoming a popular way to connect mobile computers to the Internet and other networks. The bandwidth of wireless links will probably always be limited due to properties of the physical medium and regulatory limits on the use of frequencies for radio communication. Therefore, it is necessary for network protocols to utilize the available bandwidth efficiently. Headers of IP packets are growing and the bandwidth required for transmitting headers is increasing. With the coming of I ...

Results 1 - 20 of 90

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)